

In the Claims:

Please amend the claims to read as follows (in clean format, with a marked-up version provided in an Exhibit at the end of this paper):

1. (original) A system for processing information units in a distributed processing environment where some but not all of the information units are to be sent to a plurality of different addresses, the system comprising:

a substrate including a plurality of processors, each processor for independently receiving an input information unit and creating an output information unit;

a dispatch unit carried on the substrate and coupled to the plurality of processors for receiving an input information unit and transferring the input information unit to one of said plurality of processors;

a completion unit for receiving the output information units from the plurality of processors and transmitting the output information units in an order corresponding to the input information units were received by the dispatch unit; and

a system for allowing the processing units to generate a plurality of output messages from a single input information unit, the system including an indicator created by the processing unit for each output information unit indicating whether the output information unit is the last output information unit created for the single input information unit, with the completion unit responding to the indicator and processing each message from a single input information unit until it processes an output information unit indicating that it is the last output information unit for the input information unit.

2. (original) A system of the type described in Claim 1 wherein the system for indicating whether the output information unit is the last output information unit for a single input information unit also includes a system for putting a sequence number on each output information unit created from a single input information unit.

3. (original) A system of the type described in Claim 1 further including a system for preventing the output of one processor from being passed from the completion unit until previously received messages from the same data flow have been processed, whereby later information units from the same data flow are not processed from the completion unit before earlier information units from the same data flow.

4. (original) A system of the type described in Claim 2 wherein the sequence number for each multicast message is incremented each time a new destination is provided.

5. (currently amended) A method of generating multiple output message for a single input message in a distributed information processing system, the steps of the method comprising:

providing N processing units where $N > 1$;

assigning each input information unit to one of the N processing units for processing;

processing each input information unit to create an output information unit at the one of the N processing units, with the processing including determining whether the input information unit is to be a unicast message addressed only to one recipient or if the input information unit is to be multicast to a plurality of recipients;

providing a multicast indicator indicating whether the output information unit is a part of a multicast message, and, if the multicast indicator indicates that the output information unit is a part of a multicast message, whether the output information unit is the last multicast message for a given input information unit; and

[continuing to process each message from a single input information unit until receiving the indication that the output information unit is the last multicast message from a given input information unit..]

6. (currently amended) A method for multicasting messages in a distributed information processing system, the steps of the method comprising:

associating a label field with each message based on the dataflow and adding blank bits to the label field to allow for uniquely identifying each multicast message from the same dataflow;

determining whether a message is a multicast message;

when a multicast message is determined, assigning subsequent occurrence of that message to the same processor;

incrementing the blank bits portion of the label field to distinguish each multicast message from other messages from the same multicast message; and

[continuing to process a multicast message from a single message until receiving an indicator that it is the last output from the multicast message.]
